

TITLE OF THE INVENTION
METHOD AND SYSTEM FOR PERFORMING REAL-TIME OPERATION
USING PROCESSORS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 This application is based upon and claims the
benefit of priority from prior Japanese Patent
Application No. 2003-184975, filed June 27, 2003, the
entire contents of which are incorporated herein by
reference.

10 BACKGROUND OF THE INVENTION

1. Field of the Invention

 The present invention relates to a method and
system for performing a real-time operation using
processors. More specifically, the invention relates
15 to a scheduling method in which threads for performing
a real-time operation are assigned to a plurality of
processors and a real-time processing system to which
the scheduling method is applied.

2. Description of the Related Art

20 Conventionally, computer systems such as server
computers have utilized system architecture such as a
multiprocessor and a parallel processor in order to
improve in throughput. Both of the multiprocessor and
parallel processor achieve a parallel computing
25 operation using a plurality of processing units.

 Jpn. Pat. Appln. KOKAI Publication No. 10-143380
discloses a computer system having a plurality of

processing units. This computer system includes a
single high-speed CPU, a plurality of low-speed CPUs
and a shared memory. Processes are assigned to the
high-speed and low-speed CPUs in consideration of
5 parallelism and execution time of each process.

Jpn. Pat. Appln. KOKAI Publication No. 8-180025
discloses a scheduling technique of scheduling threads
such that threads belonging to the same process are
executed on the same processor.

10 Not only the computer system but also an embedded
device, which needs to process a large amount of data
such as AV (audio video) data in real time, has
recently required that system architecture such as a
multiprocessor and a parallel processor be introduced
15 to improve in throughput.

Under the present circumstances, however, a real-
time processing system that is predicated on the above
system architecture is hardly reported.

In a real-time processing system, each operation
20 needs performing under given timing constraint. If,
however, system architecture such as a multiprocessor
and a parallel processor is applied to the real-time
processing system, latency associated with communi-
cations between threads executed by different
25 processing units becomes a serious problem.

Usually, communications for transferring data
between threads executed by different processing units

need carrying out through a buffer on a shared memory.
This is because a defect has to be prevented from
occurring in data transferred between threads even
though the threads are executed by any processor with
any timing.

BRIEF SUMMARY OF THE INVENTION

An object of the present invention is to provide
a scheduling method and real-time processing system
capable of efficiently performing communication between
threads.

According to an embodiment of the present
invention, there is provided a method of assigning a
plurality of threads to a plurality of processors, each
of the threads being a unit of execution of a real-time
operation, the method comprising selecting a tightly
coupled thread group from among the threads based on
coupling attribute information indicative of a coupling
attribute between the threads, the tightly coupled
thread group including a set of tightly coupled threads
running in cooperation with each other, and performing
a scheduling operation of dispatching the tightly
coupled threads to several of the processors that are
equal in number to the tightly coupled threads to
simultaneously execute the tightly coupled threads by
the several of the processors.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

FIG. 1 is a block diagram showing an example of a

computer system that configures a real-time processing system according to an embodiment of the present invention.

FIG. 2 is a block diagram of an MPU (master processing unit) and VPUs (versatile processing units) provided in the real-time processing system according to the embodiment of the present invention.

FIG. 3 is a diagram showing an example of a virtual address translation mechanism used in the real-time processing system according to the embodiment of the present invention.

FIG. 4 is a diagram showing an example of data mapped in real address space in the real-time processing system according to the embodiment of the present invention.

FIG. 5 is an illustration of effective address space, virtual address space and real address space in the real-time processing system according to the embodiment of the present invention.

FIG. 6 is a block diagram of a receiver for digital TV broadcast.

FIG. 7 is a diagram showing an example of a program module executed by the real-time processing system according to the embodiment of the present invention.

FIG. 8 is a table showing an example of a structural description included in the program module

shown in FIG. 7.

FIG. 9 is a chart showing a flow of data among programs corresponding to the program module shown in FIG. 7.

5 FIG. 10 is a chart showing a parallel operation of the program module shown in FIG. 7, which is performed by two VPUs.

FIG. 11 is a chart showing a pipeline operation of the program module shown in FIG. 7, which is performed
10 by two VPUs.

FIG. 12 is a diagram showing an example of an operating system in the real-time processing system according to the embodiment of the present invention.

FIG. 13 is a diagram showing another example of the operating system in the real-time processing system according to the embodiment of the present invention.
15

FIG. 14 is a diagram showing a relationship between a virtual machine OS and a guest OS in the real-time processing system according to the embodiment
20 of the present invention.

FIG. 15 is a chart showing resources that are time-divisionally assigned to a plurality of guest OSes in the real-time processing system according to the embodiment of the present invention.

25 FIG. 16 is a chart showing resources that are occupied by a specific guest OS in the real-time processing system according to the embodiment of the

present invention.

FIG. 17 is a diagram of VPU runtime environment used as a scheduler in the real-time processing system according to the embodiment of the present invention.

5 FIG. 18 is a diagram showing an example of VPU runtime environment that is implemented in the virtual machine OS used in the real-time processing system according to the embodiment of the present invention.

10 FIG. 19 is a diagram showing an example of VPU runtime environment that is implemented as a guest OS used in the real-time processing system according to the embodiment of the present invention.

15 FIG. 20 is a diagram showing an example of VPU runtime environment that is implemented in each of the guest OSes used in the real-time processing system according to the embodiment of the present invention.

20 FIG. 21 is a diagram showing an example of VPU runtime environment that is implemented in one guest OS used in the real-time processing system according to the embodiment of the present invention.

FIG. 22 is an illustration of MPU-side VPU runtime environment and VPU-side VPU runtime environment used in the real-time processing system according to the embodiment of the present invention.

25 FIG. 23 is a flowchart showing a procedure performed by the VPU-side VPU runtime environment used in the real-time processing system according to the

embodiment of the present invention.

FIG. 24 is a flowchart showing a procedure performed by the MPU-side VPU runtime environment used in the real-time processing system according to the
5 embodiment of the present invention.

FIG. 25 is an illustration of threads belonging to a tightly coupled thread group and executed by different processors in the real-time processing system according to the embodiment of the present invention.

10 FIG. 26 is an illustration of interaction between tightly coupled threads in the real-time processing system according to the embodiment of the present invention.

FIG. 27 is an illustration of mapping of local
15 storages of VPU's executing partner threads in effective address spaces of the tightly coupled threads in the real-time processing system according to the embodiment of the present invention.

FIG. 28 is an illustration of allocation of
20 processors to threads belonging to a loosely coupled thread group in the real-time processing system according to the embodiment of the present invention.

FIG. 29 is an illustration of interaction between
loosely coupled threads in the real-time processing
25 system according to the embodiment of the present invention.

FIG. 30 is an illustration of a relationship

between processes and threads in the real-time processing system according to the embodiment of the present invention.

FIG. 31 is a flowchart showing a procedure for performing a scheduling operation in the real-time processing system according to the embodiment of the present invention.

FIG. 32 is an illustration of a first issue of mapping of local storages in the real-time processing system according to the embodiment of the present invention.

FIG. 33 is an illustration of a relationship between a physical VPU and a logical VPU in the real-time processing system according to the embodiment of the present invention.

FIG. 34 is an illustration of a second issue of mapping of local storages in the real-time processing system according to the embodiment of the present invention.

FIG. 35 is an illustration of a shared model of effective address space in the real-time processing system according to the embodiment of the present invention.

FIG. 36 is an illustration of a shared model of virtual address space in the real-time processing system according to the embodiment of the present invention.

FIG. 37 is an illustration of an unshared model in the real-time processing system according to the embodiment of the present invention.

5 FIG. 38 is a first diagram describing a change in mapping of local storages in the real-time processing system according to the embodiment of the present invention.

10 FIG. 39 is a second diagram describing a change in mapping of local storages in the real-time processing system according to the embodiment of the present invention.

15 FIG. 40 is a third diagram describing a change in mapping of local storages in the real-time processing system according to the embodiment of the present invention.

FIG. 41 is a fourth diagram describing a change in mapping of local storages in the real-time processing system according to the embodiment of the present invention.

20 FIG. 42 is a fifth diagram describing a change in mapping of local storages in the real-time processing system according to the embodiment of the present invention.

25 FIG. 43 is a flowchart showing a procedure for address administration performed to change the mapping of local storages in the real-time processing system according to the embodiment of the present invention.

FIG. 44 is an illustration of a change in mapping between a memory and local storages in the real-time processing system according to the embodiment of the present invention.

5 FIG. 45 is a flowchart showing a procedure for the change in mapping between the memory and local storages in the real-time processing system according to the embodiment of the present invention.

10 FIG. 46 is a diagram showing a state transition of threads in the real-time processing system according to the embodiment of the present invention.

15 FIG. 47 is a chart illustrating a relationship between a thread and execution terms in the real-time processing system according to the embodiment of the present invention.

FIG. 48 is a chart of tightly coupled threads running at once in an execution term in the real-time processing system according to the embodiment of the present invention.

20 FIG. 49 is a chart showing a periodic execution model in the real-time processing system according to the embodiment of the present invention.

25 FIG. 50 is a chart showing an aperiodic execution model in the real-time processing system according to the embodiment of the present invention.

FIG. 51 is an illustration of a task graph.

FIG. 52 is an illustration of the principle of a

reservation graph used in the real-time processing system according to the embodiment of the present invention.

5 FIG. 53 is an illustration of an example of a reservation graph used in the real-time processing system according to the embodiment of the present invention.

10 FIG. 54 is a diagram illustrating a hierarchical scheduler used in the real-time processing system according to the embodiment of the present invention.

FIG. 55 is a chart illustrating examples of parameters used for scheduling in the hard real-time class by the real-time processing system according to the embodiment of the present invention.

15 FIG. 56 is an illustration of absolute timing constraint used in the real-time processing system according to the embodiment of the present invention.

20 FIG. 57 is an illustration of relative timing constraint used in the real-time processing system according to the embodiment of the present invention.

FIG. 58 is an illustration of mutual exclusive constraint used in the real-time processing system according to the embodiment of the present invention.

25 FIG. 59 is a table illustrating synchronization mechanisms in the real-time processing system according to the embodiment of the present invention.

FIG. 60 is a flowchart showing a procedure for

selectively using the synchronization mechanisms in the real-time processing system according to the embodiment of the present invention.

5 FIG. 61 is a diagram showing an example of a reservation graph used in the real-time processing system according to the embodiment of the present invention.

10 FIG. 62 is a diagram showing an example of a reservation request created in the real-time processing system according to the embodiment of the present invention.

15 FIG. 63 is a chart showing an example of scheduling performed by the real-time processing system according to the embodiment of the present invention on the basis of the reservation request shown in FIG. 62.

 FIG. 64 is a chart illustrating a first example of scheduling of software pipeline type performed by the real-time processing system according to the embodiment of the present invention.

20 FIG. 65 is a chart illustrating a second example of scheduling of software pipeline type performed by the real-time processing system according to the embodiment of the present invention.

25 FIG. 66 is a chart illustrating a first example of scheduling performed in consideration of the number of buffers by the real-time processing system according to the embodiment of the present invention.

FIG. 67 is a chart illustrating a second example of scheduling performed in consideration of the number of buffers by the real-time processing system according to the embodiment of the present invention.

5 FIG. 68 is a chart illustrating a third example of scheduling performed in consideration of the number of buffers by the real-time processing system according to the embodiment of the present invention.

10 FIG. 69 is a flowchart of procedures for the scheduling performed in consideration of the number of buffers by the real-time processing system according to the embodiment of the present invention.

15 FIG. 70 is a diagram showing an example of a reservation graph having a hierarchical structure used in the real-time processing system according to the embodiment of the present invention.

20 FIG. 71 is a diagram showing an example of a reservation request which is created by the real-time processing system according to the embodiment of the present invention and which takes into consideration of the tightly coupled thread group.

25 FIG. 72 is a chart showing an example of scheduling performed by the real-time processing system according to the embodiment of the present invention on the basis of the reservation request shown in FIG. 71.

 FIG. 73 is a diagram showing an example of a reservation list used in the real-time processing

system according to the embodiment of the present invention.

FIG. 74 is a flowchart showing a procedure for reserving an execution term in the real-time processing system according to the embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

An embodiment of the present invention will now be described with reference to the accompanying drawings.

FIG. 1 shows an example of a configuration of a computer system for achieving a real-time processing system according to an embodiment of the present invention. The computer system is an information processing system that performs various operations, which need to be done in real time, under timing constraint. The computer system can be used as not only a general-purpose computer but also an embedded system for various electronic devices to perform operations that need to be done in real time.

Referring to FIG. 1, the computer system comprises an MPU (master processing unit) 11, a plurality of VPUs (versatile processing units) 12, a connecting device 13, a main memory 14 and an I/O (input/output) controller 15. The MPU 11, VPUs 12, main memory 14 and IO controller 15 are connected to each other by the connecting device 13. The connecting device 13 is formed of a bus or an inter-connection network such as

a crossbar switch. If a bus is used for the connecting device 13, it also can be shaped like a ring. The MPU 11 is a main processor that controls an operation of the computer system. The MPU 11 mainly executes an OS
5 (operating system). Some functions of the OS can be executed by the VPUs 12 and IO controller 15. Each of the VPUs 12 is a processor for performing various operations under the control of the MPU 11. The MPU 11 distributes the operations (tasks) to be performed to
10 the VPUs 12 in order to perform these operations (tasks) in parallel. The operations can thus be performed at high speed and with high efficiency. The main memory 14 is a main storage device (shared memory) that is shared by the MPU 11, VPUs 12 and I/O
15 controller 15. The main memory 14 stores the OS and application programs. The I/O controller 15 is connected to one or more I/O devices 16. The controller 15 is also referred to as a bridge device.

The connecting device 13 has a QoS (quality of
20 service) function that guarantees a data transfer rate. The QoS function is fulfilled by transferring data through the connecting device 13 at a reserved bandwidth (transfer rate). The QoS function is used when write data needs transmitting to the memory 14
25 from one VPU 12 at e.g., 5 Mbps or when data needs transferring between one VPU 12 and another VPU 12 at e.g., 100 Mbps. Each of the VPUs 12 designates

(reserves) a bandwidth (transfer rate) for the connecting device 13. The connecting device 13 assigns the designated bandwidth to the VPU 12 by priority. If a bandwidth is reserved for data transfer of a VPU 12, it is secured even though another VPU 12, MPU 11 or IO controller 15 transfers a large amount of data during the data transfer of the former VPU 12. The QoS function is particularly important to computers that perform real-time operations.

The computer system shown in FIG. 1 comprises one MPU 11, four VPUs 12, one memory 14 and one IO controller 15. The number of VPUs 12 is not limited. The system can be configured without MPU and, in this case, one VPU 12 performs the operation of the MPU 11. In other words, one VPU 12 serves as a virtual MPU 11.

FIG. 2 shows an MPU 11 and VPUs 12. The MPU 11 includes a processing unit 21 and a memory management unit 22. The processing unit 21 accesses the memory 14 through the memory management unit 22. The memory management unit 22 performs a virtual memory management function and also manages a cache memory in the memory management unit 22. Each of the VPUs 12 includes a processing unit 31, a local storage (local memory) 32 and a memory controller 33. The processing unit 31 of each VPU 12 can gain direct access to the local storage 32 in the same VPU 12. The memory controller 33 serves as a DMA (direct memory access) controller that

transfers data between the local storage 32 and memory 14. The memory controller 33 is so configured to utilize the QoS function of the connecting device 13 and has a function of designating a bandwidth and that of inputting/outputting data at the designated bandwidth. The memory controller 33 also has the same virtual memory management function as that of the memory management unit 22 of the MPU 11. The processing unit 31 uses the local storage 32 as a main memory. The processing unit 31 does not gain direct access to the memory 14 but instructs the memory controller 33 to transfer the contents of the memory 14 to the local storage 32. The processing unit 31 accesses the local storage 32 to read/write data. Moreover, the processing unit 31 instructs the memory controller 33 to write the contents of the local storage 32 to the memory 14.

The memory management unit 22 of the MPU 11 and the memory controllers 33 of the VPUs 12 perform virtual memory management as shown in FIG. 3. The address viewed from the processing unit 21 of the MPU 11 or the memory controllers 33 of the VPUs 12 is a 64-bit address as indicated in the upper part of FIG. 3. In the 64-bit address, an upper 36-bit portion indicates a segment number, a middle 16-bit portion indicates a page number, and a lower 12-bit portion indicates a page offset. The memory management unit 22

and memory controllers 33 each include a segment table 50 and a page table 60. The segment table 50 and page table 60 convert the 64-bit address into the real address space that is actually accessed through the connecting device 13.

For example, the following data items are mapped in the real address (RA) space viewed from the MPU 11 and each VPU 12, as shown in FIG. 4.

1. Memory 14 (main storage device)
2. Control registers of MPU 11
3. Control registers of VPUs 12
4. Local storages of VPUs 12
5. Control registers of I/O devices (including control registers of I/O controller 15)

The MPU 11 and VPUs 12 can access any address in the real address space by the virtual memory management function in order to read/write data items 1 to 5 described above. It is particularly important to be able to access the real address space and thus access the local storage 32 of any VPU 12 from the MPU 11 and VPUs 12 and even from the I/O controller 15. Furthermore, the segment table 50 or page table 60 can prevent the contents of the local storage 32 of each VPU 12 from being read or written freely.

FIG. 5 shows memory address spaces managed by the virtual memory management function shown in FIG. 3. It is the EA (effective address) space that is viewed

directly from the programs executed on the MPU 11 or
VPUs 12. An effective address is mapped in the VA
(virtual address) space by the segment table 50. A
virtual address is mapped in the RA (real address)
5 space by the page table 60. The RA space has a
structure as shown in FIG. 4.

The MPU 11 can manage the VPUs 12 using a hardware
mechanism such as a control register. For example, the
MPU 11 can read/write data from/to the register of each
10 VPU 12 and start/stop each VPU 12 to execute programs.
Communication and synchronization between the MPU 11
and each of the VPUs 12 can be performed by means of a
hardware mechanism such as a mailbox and an event flag,
as can be communication and synchronization between the
15 VPUs 12.

The computer system according to the present
embodiment allows an operation of an electric device,
which makes a stringent demand on real-time operations
as conventionally implemented by hardware, to be
20 carried out by software. For example, one VPU 12
performs a computation corresponding to some hardware
components that compose the electric device and
concurrently another VPU 12 performs a computation
corresponding to other hardware components that compose
25 the electric device.

FIG. 6 simply shows a hardware structure of a
receiver for digital TV broadcast. In this receiver, a

DEMUX (demultiplexer) circuit 101 divides a received broadcast signal into compressing-encoded data streams corresponding to audio data, video data and subtitle data. An A-DEC (audio decoder) circuit 102 decodes the
5 compressing-encoded audio data stream. A V-DEC (video decoder) circuit 103 decodes the compressing-encoded video data stream. The decoded video data stream is sent to a PROG (progressive conversion) circuit 105 and converted into a progressive video signal. The
10 progressive video signal is sent to a BLEND (image blending) circuit 106. A TEXT (subtitle data processing) circuit 104 converts the compressing-encoded subtitle data stream into a subtitle video signal and sends it to the BLEND circuit 106. The
15 BLEND circuit 106 blends the video signal sent from the PROG circuit 105 and the subtitle video signal sent from the TEXT circuit 104 and outputs the blended signal as a video stream. A series of operations as described above is repeated at a video frame rate
20 (e.g., 30, 32 or 60 frames per second).

In order to perform operations of the hardware shown in FIG. 6 by software, the present embodiment provides a program module 100 as shown in FIG. 7. The program module 100 is an application program for
25 causing the computer system to perform the operations of the DEMUX circuit 101, A-DEC circuit 102, V-DEC circuit 103, TEXT circuit 104, PROG circuit 105 and

BLEND circuit 106 shown in FIG. 6. The application program is described by multi-thread programming, and is structured as a group of threads for executing a real-time operation. The real-time operation includes a combination of a plurality of tasks. The program module 100 contains a plurality of programs (a plurality of routines) each executed as a thread. Specifically, the program module 100 contains a DEMUX program 111, an A-DEC program 112, a V-DEC program 113, a TEXT program 114, a PROG program 115 and a BLEND program 116. These programs 111 to 116 are programs describing procedures of tasks corresponding to operations (DMUX operation, A-DEC operation, V-DEC operation, TEXT operation, PROG operation, BLEND operation) of the circuits 101 to 106. More specifically, when the program module 100 runs, a thread corresponding to each of the programs 111 to 116 is generated, and dispatched to one or more VPUs 12 and executed thereon. A program corresponding to the thread dispatched to the VPU 12 is loaded to the local storage 32 of the VPU 12, and the thread executes the program on the local storage 32. The program module 100 is obtained by packaging the programs 111 to 116, which correspond to hardware modules for configuring a receiver for digital TV broadcast, with data called a structural description 117.

The structural description 117 is information

indicative of how the programs (threads) in the program module 100 are combined and executed. The structural description 117 includes information indicative of a relationship in input/output between the programs 111 to 116 and costs (time) necessary for executing each of the programs 111 to 116. FIG. 8 shows an example of the structural description 117.

The structural description 117 shows modules (programs in the program module 100) each executed as a thread and their corresponding inputs, outputs, execution costs, and buffer sizes necessary for the outputs. For example, the V-DEC program of No. (3) receives the output of the DEMUX program of No. (1) as an input and transmits its output to the PROG program of No. (5). The buffer necessary for the output of the V-DEC program is 1 MB and the cost for executing the V-DEC program in itself is 50. The cost can be described in units of time (time period) necessary for executing the program, or step number of the program. It also can be described in units of time required for executing the program by a virtual processor having some virtual specifications. Since the VPU specifications and performance may vary from computer to computer, it is desirable to describe the cost in such virtual units. If the programs are executed according to the structural description 117 shown in FIG. 8, data flows among the programs as illustrated in FIG. 9.

The structural description 117 also shows coupling attribute information, which indicates a coupling attribute between threads corresponding to the programs 111 to 116, as thread parameters. The coupling attribute includes two different attributes of a tightly coupled attribute and a loosely coupled attribute. A plurality of threads having the tightly coupled attribute are executed in cooperation with each other and referred to as a tightly coupled thread group. The computer system of the present embodiment schedules the threads belonging to each tightly coupled thread group such that the threads belonging to the same tightly coupled thread group can simultaneously be executed by different VPUs. A plurality of threads having the loosely coupled attribute is referred to as a loosely coupled thread group. A programmer can designate a coupling attribute between threads corresponding to the programs 11 to 16 using thread parameters. The tightly and loosely coupled thread groups will be described in detail with reference to FIG. 25 et seq. The thread parameters including the coupling attribute information can be described directly as codes in the programs 111 to 116, not as the structural description 117.

Referring to FIGS. 10 and 11, there now follows descriptions as to how the computer system of the present embodiment executes the programs 111 to 116.

Assume here that the computer system includes two VPUs of VPU0 and VPU1. FIG. 10 shows time for assigning the programs to each of the VPUs when video data of 30 frames is displayed per second. Audio and video data for one frame is output in one period (1/30 second). First, the VPU0 executes the DEMUX program to perform the DEMUX operation and writes its resultant audio, video and subtitle data to the buffers. After that, the VPU1 executes the A-DEC program and TEXT program to perform the A-DEC operation and the TEXT operation in sequence and writes their results to the buffers. Then, the VPU0 executes the V-DEC program to perform the V-DEC operation and writes its result to the buffer. The VPU0 executes the PROG program to perform the PROG operation and writes its result to the buffer. Since the VPU1 has already completed the TEXT program at this time, the VPU0 executes the last BLEND program to perform the BLEND operation, in order to create final video data. The above processing is repeated for every period.

An operation to determine which program is executed by each of the VPUs and when it is done to perform a desired operation without delay is called scheduling. A module to carry out the scheduling is called a scheduler. In the present embodiment, the scheduling is carried out based on the above structural description 117 contained in the program module 100.

FIG. 11 shows the programs executed when video data of 60 frames is displayed per second. FIG. 11 differs from FIG. 10 as follows. In FIG. 11, data of 60 frames needs to be processed per second, whereas in FIG. 10, data of 30 frames is processed per second and thus data processing for one frame can be completed in one period ($1/30$ second). In other words, one-frame data processing cannot be completed in one period ($1/60$ second) and thus a software pipeline operation that spans a plurality of (two) periods is performed in FIG. 11. For example, in period 1, the VPU0 executes the DEMUX program and V-DEC program for the input signal. After that, in period 2, the VPU1 executes the A-DEC, TEXT, PROG and BLEND programs and outputs final video data. In period 2, the VPU0 executes the DEMUX and V-DEC programs in the next frame. The DEMUX and V-DEC programs of the VPU0 and the A-DEC, TEXT, PROG and BLEND programs of the VPU1 are executed over two periods as a pipeline operation.

The program module 100 shown in FIG. 7 can be recorded in advance in a flash ROM and a hard disk in a device incorporating the computer system of the present embodiment, or circulated through a network. In this case, the contents of operations to be performed by the computer system vary according to the type of a program module downloaded through the network. Thus, the device incorporating the computer system can perform

the real-time operation corresponding to each of various pieces of dedicated hardware. If new player software, decoder software and encryption software necessary for reproducing new contents are distributed together with the contents as program modules executable by the computer system, any device incorporating the computer system can reproduce the contents within acceptable limits of ability.

Operating System

When only one OS (operating system) 201 is loaded into the computer system of the present embodiment, it manages all real resources (MPU 11, VPUs 12, memory 14, I/O controller 15, I/O device 16, etc.), as shown in FIG. 12.

On the other hand, a plurality of Oses can be performed at once using a virtual machine system. In this case, as shown in FIG. 13, a virtual machine OS 301 is loaded into the computer system to manage all real resources (MPU 11, VPUs 12, memory 14, I/O controller 15, I/O device 16, etc.). The virtual machine OS 301 is also referred to as a host OS. One or more Oses 302 and 303, which are also referred to as guest Oses, are loaded on the virtual machine OS 301. Referring to FIG. 14, the guest Oses 302 and 303 each run on a computer including virtual machine resources given by the virtual machine OS 301 and provide various services to application programs managed by the guest

OSes 302 and 303. In the example of FIG. 14, the guest OS 302 appears as if it operated on a computer including one MPU 11, two VPUs 12 and one memory 14, and the guest OS 303 appears as if it operated on a computer including one MPU 11, four VPUs 12 and one memory 14. The virtual machine OS 301 manages which one of VPUs 12 of the real resources actually corresponds to a VPU 12 viewed from the guest OS 302 and a VPU 12 viewed from the guest OS 303. The guest OSes 302 and 303 need not be aware of the correspondence.

The virtual machine OS 301 schedules the guest OSes 302 and 303 to allocate all the resources in the computer system to the guest OSes 302 and 303 on a time-division basis. Assume that the guest OS 302 carries out a real-time operation. To perform the operation thirty times per second at an exact pace, the guest OS 302 sets its parameters to the virtual machine OS 301. The virtual machine OS 301 schedules the guest OS 302 to reliably assign necessary operation time to the guest OS 302 once per 1/30 second. The operation time is assigned to a guest OS that does not require a real-time operation by priority lower than a guest OS that requires a real-time operation. FIG. 15 shows that the guest OSes 302 and 303 run alternately, representing time by the horizontal axis. While the guest OS 302 (OS1) is running, the MPU 11 and all the

VPUs 12 are used as resources of the guest OS 302 (OS1). While the guest OS 303 (OS2) is running, the MPU 11 and all the VPUs 12 are used as resources of the guest OS 303 (OS2).

5 FIG. 16 shows an operation mode different from that in FIG. 15. There is a case where it is to be wished that a VPU 12 be used continuously according to target applications. This case corresponds to, for example, an application that necessitates continuing to
10 monitor data and events all the time. The scheduler of the virtual machine OS 301 manages the schedule of a specific guest OS such that the guest OS occupies a specific VPU 12. In FIG. 16, a VPU 3 is designated as a resource exclusively for a guest OS 302 (OS1). Even
15 though the virtual machine OS 301 switches the guest OS 302 (OS1) and guest OS 303 (OS2) to each other, the VPU 3 always continues to operate under the control of the guest OS 302 (OS1).

 In order to execute programs using a plurality of
20 VPUs 12 in the present embodiment, a software module called a VPU runtime environment is used. The soft module includes a scheduler for scheduling threads to be assigned to the VPUs 12. When only one OS 201 is implemented on the computer system of the present
25 embodiment, a VPU runtime environment 401 is implemented on the OS 201 as illustrated in FIG. 17. The VPU runtime environment 401 can be implemented in

the kernel of the OS 201 or in a user program. It can also be divided into two for the kernel and user program to run in cooperation with each other. When one or more guest OSes run on the virtual machine OS 301, the following modes are provided to implement the VPU runtime environment 401:

1. Mode of implementing the VPU runtime environment 401 in the virtual machine OS 301 (FIG. 18).

2. Mode of implementing the VPU runtime environment 401 as one OS managed by the virtual machine OS 301 (FIG. 19). In FIG. 19, the guest OS 304 running on the virtual machine OS 301 is the VPU runtime environment 401.

3. Mode of implementing a dedicated VPU runtime environment in each of the guest OSes managed by the virtual machine OS 301 (FIG. 20). In FIG. 20, the VPU runtime environments 401 and 402 are implemented in their respective guest OSes 302 and 303. The VPU runtime environments 401 and 402 run in association with each other, if necessary, using a function of communication between the guest OSes provided by the virtual machine OS 301.

4. Mode of implementing the VPU runtime environment 401 in one of the guest OSes managed by the virtual machine OS 301 (FIG. 21). A guest OS 303 having no VPU runtime environment utilizes the VPU

runtime environment 401 of a guest OS 302 using a function of communication between the guest OSes provided by the virtual machine OS 301.

The above modes have the following merits:

5 Merits of Mode 1

 The scheduling of a guest OS managed by the virtual machine OS 301 and that of the VPU can be combined into one. Thus, the scheduling can be done efficiently and finely and the resources can be used
10 effectively; and

 Since the VPU runtime environment can be shared among a plurality of guest OSes, a new VPU runtime environment need not be created when a new guest OS is introduced.

15 Merits of Mode 2

 Since a scheduler for the VPUs can be shared among guest OSes on the virtual machine OS, the scheduling can be performed efficiently and finely and the resources can be used effectively;

20 Since the VPU runtime environment can be shared among a plurality of guest OSes, a new VPU runtime environment need not be created when a new guest OS is introduced; and

 Since the VPU runtime environment can be created
25 without depending upon the virtual machine OS or a specific guest OS, it can be standardized easily and replaced with another. If a VPU runtime environment

suitable for a specific embedded device is created to perform scheduling utilizing the characteristics of the device, the scheduling can be done with efficiency.

Merit of Mode 3

5 Since the VPU runtime environment can optimally be implemented in each guest OS, the scheduling can be performed efficiently and finely and the resources can be used effectively.

Merit of Mode 4

10 Since the VPU runtime environment need not be implemented in all the guest OSes, a new guest OS is easy to add.

As is evident from the above, all the modes 1 to 4 can be used to implement the VPU runtime environment.

15 Any other modes can be used when the need arises.

Service Provider

20 In the computer system according to the present embodiment, the VPU runtime environment 401 provides various services (a communication function using a network, a function of inputting/outputting files, calling a library function such as a codec, interfacing with a user, an input/output operation using an I/O device, reading of date and time, etc.) as well as functions of managing and scheduling various resources
25 (operation time of each VPU, a memory, bandwidth of a connection device, etc.) associated with the VPUs 12. These services are called from application programs

running on the VPUs 12. If a simple service is called,
it is processed by service programs on the VPUs 12. A
service that cannot be processed only by the VPUs 12,
such as communication processing and file processing,
5 is processed by service programs on the MPU 11. The
programs that provide such services are referred to as
a service provider (SP).

FIG. 22 shows one example of the VPU runtime
environment. The principal part of the VPU runtime
10 environment is present on the MPU 11 and corresponds to
an MPU-side VPU runtime environment 501. A VPU-side
VPU runtime environment 502 is present on each of the
VPUs 12 and has only the minimum function of carrying
out a service that can be processed in the VPU 12. The
15 function of the MPU-side VPU runtime environment 501 is
roughly divided into a VPU controller 511 and a service
broker 512. The VPU controller 511 chiefly provides
a management mechanism, a synchronization mechanism,
a security management mechanism and a scheduling
20 mechanism for various resources (operation time of
each VPU, a memory, a virtual space, bandwidth of a
connection device, etc.) associated with the VPUs 12.
It is the VPU controller 511 that dispatches programs
to the VPUs 12 based on the results of scheduling.
25 Upon receiving a service request called by the
application program on each VPU 12, the service broker
512 calls an appropriate service program (service

provider) and provides the service.

Upon receiving a service request called by the application program on each VPU 12, the VPU-side VPU runtime environment 502 processes only services that are processable in the VPU 12 and requests the service broker 512 to process services that are not processable therein.

FIG. 23 shows a procedure for processing a service request by the VPU-side VPU runtime environment 502.

10 Upon receiving a service call from an application program (step S101), the VPU-side VPU runtime environment 502 determines whether the service can be processed therein (step S102). If the service can be processed, the VPU runtime environment 502 executes the service and returns its result to the calling part (steps S103 and S107). If not, the VPU runtime environment 502 determines whether a service program that can execute the service is registered as one executable on each VPU 12 (step S104). If the service program is registered, the VPU runtime environment 502 executes the service program and returns its result to the calling part (steps S105 and S107). If not, the VPU runtime environment 502 requests the service broker 512 to execute the service program and returns a result of the service from the service broker 512 to the calling part (steps S106 and S107).

FIG. 24 shows a procedure for processing a

service, which is requested by the VPU-side VPU runtime environment 502, by the service broker 512 of the MPU-side VPU runtime environment 501. Upon receiving a service call from the VPU-side VPU runtime environment 502 (step S111), the service broker 512 determines whether the service can be processed in the VPU runtime environment 501 (step S112). If the service can be processed, the service broker 512 executes the service and returns its result to the VPU-side VPU runtime environment 502 of the calling part (steps S113 and S114). If not, the service broker 512 determines whether a service program that can execute the service is registered as one executable on the MPU 11 (step S114). If the service program is registered, the service broker 512 executes the service program and returns its result to the VPU-side VPU runtime environment 502 of the calling part (steps S116 and S114). If not, the service broker 512 returns an error to the VPU-side VPU runtime environment 502 of the calling part (step S117).

Results reply to some service requests issued from the program to be executed by each VPU 12, and no results reply to other service requests. The destination of the reply is usually a thread that issues a service request; however, another thread, a thread group or a process can be designated as the destination of the reply. It is thus favorable that

the destination be included in a message to request a service. The service broker 512 can be realized using a widely used object request broker.

Real-time Operation

5 The computer system according to the present embodiment serves as a real-time processing system. The operations to be performed by the real-time processing system are roughly divided into the following three types:

- 10 1. Hard real-time operation
 2. Soft real-time operation
 3. Best effort operation (non-real-time operation)

The hard and soft real-time operations are a so-called real-time operation. The real-time processing system
15 of the present embodiment has concepts of both thread and process like a number of existing OSes. First, the thread and process in the real-time processing system will be described.

The thread has the following three classes:

- 20 1. Hard real-time class

Timing requirements are very important. This thread class is used for such an important application as to cause a grave condition when the requirements are not met.

- 25 2. Soft real-time class

This thread class is used for an application whose quality simply lowers even if the timing requirements

are not met.

3. Best effort class

This thread class is used for an application including no timing requirements.

5 In the present embodiment, the thread is a unit of execution for the real-time operation. The threads have their related programs that are to be executed by the threads. Each of the threads holds its inherent information that is called a thread context. The
10 thread context contains, for example, information of a stack and values stored in the register of the processor.

 In the real-time processing system, there are two different threads of MPU and VPU threads. These two
15 threads are classified by processors (MPU 11 and VPU 12) that execute the threads and their models are identical with each other. The thread context of the VPU thread includes the contents of the local storage 32 of the VPU 12 and the conditions of a DMA controller
20 of the memory controller 33.

 A group of threads is called a thread group. The thread group has the advantage of efficiently and easily performing, e.g., an operation of giving the same attribute to the threads of the group. The thread
25 group in the hard or soft real-time class is roughly divided into a tightly coupled thread group and a loosely coupled thread group. The tightly coupled

thread group and loosely coupled thread group are discriminated from each other by attribute information (coupling attribute information) added to the thread groups. The coupling attribute of the thread groups can explicitly be designated by the codes in the application programs or the above-described structural description.

The tightly coupled thread group is a thread group that is made up of threads running in cooperation with each other. In other words, the threads belonging to the tightly coupled thread group tightly collaborate with each other. The tightly collaboration implies an interaction such as frequent communication and synchronization between threads or an interaction that decreases in latency. The threads belonging to the same tightly coupled thread group are always executed simultaneously. On the other hand, the loosely coupled thread group is a thread group that obviates a tightly collaboration between threads belonging to the group. The threads belonging to the loosely coupled thread group carry out communications for transferring data through the buffer on the memory 14.

Tightly Coupled Thread Group

As shown in FIG. 25, different VPUs are allocated to the threads of the tightly coupled thread group and the threads are executed at the same time. These threads are called tightly coupled threads. The

execution terms of the tightly coupled threads are reserved in their respective VPUs, and the tightly coupled threads are executed at the same time. In FIG. 25, a tightly coupled thread group includes two
5 tightly coupled threads A and B and the threads A and B are executed at once by the VPU0 and VPU1, respectively. The real-time processing system of the present embodiment ensures that the threads A and B are executed at once by different VPUs. One of the threads
10 can directly communicate with the other thread through a local storage or control register of the VPU that executes the other thread.

FIG. 26 illustrates communication between threads A and B, which is performed through the local storages
15 of VPU0 and VPU1 that execute the threads A and B, respectively.

In the VPU0 that executes the thread A, an RA space corresponding to the local storage 32 of the VPU1 that executes the thread B is mapped in part of an EA space of the thread A. For this mapping, an address
20 translation unit 331 provided in the memory controller 33 of the VPU0 performs address translation using a segment table and page table. The address translation unit 331 converts (translates) a part of the EA space
25 of the thread A to the RA space corresponding to the local storage 32 of the VPU1, thereby to map the RA space corresponding to the local storage 32 of the VPU1

in part of the EA space of the thread A.

In the VPU1 that executes the thread B, an RA space corresponding to the local storage 32 of the VPU0 that executes the thread A is mapped in part of an EA space of the thread B. For this mapping, an address translation unit 331 provided in the memory controller 33 of the VPU1 performs address translation using the segment table and page table. The address translation unit 331 converts a part of the EA space of the thread B to the RA space corresponding to the local storage 32 of the VPU0, thereby to map the RA space corresponding to the local storage 32 of the VPU0 in part of the EA space of the thread B.

FIG. 27 shows mapping of local storage (LS1) 32 of the VPU1 executing the thread B in the EA space of the thread A executed by the VPU0 and mapping of local storage (LS0) 32 of the VPU0 executing the thread A in the EA space of the thread B executed by the VPU1. For example, when data to be transferred to the thread B is prepared on the local storage LS0, the thread A sets a flag indicative of this preparation in the local storage LS0 of the VPU0 or the local storage LS1 of the VPU1 that executes the thread B. In response to the setting of the flag, the thread B reads the data from the local storage LS0.

According to the present embodiment described above, tightly coupled threads can be specified by the

coupling attribute information, and the tightly coupled threads A and B are sure to be executed at once by different VPUs, respectively. Thus, an interaction of communication and synchronization between the threads A and B can be performed more lightly without delay.

Loosely Coupled Thread Group

The execution term of each of threads belonging to the loosely coupled thread group depends upon the relationship in input/output between the threads. Even though the threads are subject to no constraints of execution order, it is not ensured that they are executed at the same time. The threads belonging to the loosely coupled thread group are called loosely coupled threads. FIG. 28 shows a loosely coupled thread group including two threads C and D as loosely coupled threads, which are executed by their respective VPU0 and VPU1. The threads C and D differ in execution term as is apparent from FIG. 28. Communication between the threads C and D is carried out by the buffer prepared on the main memory 14 as shown in FIG. 29. The thread C executed by the VPU0 writes data, which is prepared in the local storage LS0, to the buffer prepared on the main memory 14 by DMA transfer. The thread D executed by the VPU1 reads data from the buffer on the main memory 14 and writes it to the local storage LS1 by DMA transfer when the thread D starts to run.

Process and Thread

As shown in FIG. 30, a process includes one address space and one or more threads. The threads can be included in the process regardless of their number and type. For example, only VPU threads can be included in the process and so can be a mixture of VPU and MPU threads. As a thread holds a thread context as its inherent information, a process holds a process context as its inherent information. The process context contains both an address space inherent in the process and thread contexts of all threads included in the process. The address space can be shared among all the threads of the process. One process can include a plurality of thread groups, but one thread group cannot belong to a plurality of processes. Thus, a thread group belonging to a process is inherent in the process.

In the real-time processing system of the present embodiment, there are two models of a thread first model and an address space first model as method for creating a new thread. The address space first model is the same as that adopted in the existing OS and thus can be applied to both the MPU and VPU threads. On the other hand, the thread first model can be applied only to the VPU threads and is peculiar to the real-time processing system of the present embodiment. In the thread first model, the existing thread (which is one

for creating a new thread, i.e., a parent thread of the new thread) first designates a program to be executed by a new thread and causes the new thread to start to execute the program. The program is then stored in the local storage of the VPU and starts to run from a given address. Since no address space is related to the new thread at this time, the new thread can gain access to the local storage of the VPU and not to the memory 14. After that, when the need arises, the new thread in itself calls a service of VPU runtime environment and creates an address space. The address space is related to the new thread, and the new thread can gain access to the memory 14. In the address space first model, the existing thread creates a new address space or designates the existing address space, and arranges program, which is to execute by the new thread, in the address space. Then, the new thread starts to run the programs. The merit of the thread first model is that a thread can be executed only by the local storage to reduce overhead costs required for generating, dispatching and exiting the thread.

Scheduling of Threads

A scheduling operation performed by the VPU runtime environment 401 will now be described with reference to the flowchart shown in FIG. 31. The scheduler in the VPU runtime environment 401 checks a coupling attribute between threads based on coupling

attribute information added to each group of threads to be scheduled (step S121). The scheduler determines whether each thread group is a tightly coupled thread group or a loosely coupled thread group (step S122).

5 The coupling attribute is checked referring to the descriptions of threads in program codes or thread parameters in the above structural description 117. If the tightly and loosely coupled thread groups are each specified, the threads to be scheduled are separated
10 into the tightly and loosely coupled thread groups.

The scheduling of threads belonging to the tightly coupled thread group is performed as follows. In order to execute threads of a tightly coupled thread group, which are selected from the threads to be scheduled, by
15 their respective VPUs at once, the scheduler in the VPU runtime environment 401 reserves an execution term of each of the VPUs, whose number is equal to that of the threads, and dispatches the threads to the VPUs at once (step S123). The scheduler maps an RA space in part of
20 an EA space of a thread using the address translation unit 331 in a VPU that executes the thread (step S124), the RA space corresponding to the local storage of a VPU that executes a partner thread interacting with the former thread. As for the threads belonging to the
25 loosely coupled thread group which are selected from the threads to be scheduled, the scheduler dispatches the threads in sequence to one or more VPUs based on

the relationship in input/output between the threads
(step S125).

5 If a tightly coupled thread group, which is a set
of threads running in cooperation with each other, is
selected based on the coupling attribute information,
it can be ensured that the threads belonging to the
tightly coupled thread group are executed at once by
different processors. Consequently, communication
between threads can be achieved by a lightweight
10 mechanism of gaining direct access to, e.g., the
registers of processors that execute their partner
threads each other. The communication can thus be
performed lightly and quickly.

Mapping of Local Storage

15 In the real-time processing system of the present
embodiment, when MPU and VPU threads or VPU threads
perform an operation of communication or synchroniza-
tion in cooperation with each other, it is necessary to
access the local storage of the partner VPU thread.
20 For example, a more lightweight, high-speed synchro-
nization mechanism is implemented by a synchronous
variable assigned on the local storage. It is thus
necessary that the local storage of a VPU 12 be
accessed directly by another VPU 12 or the MPU 11. If
25 a segment table or page table is set appropriately when
the local storage of a VPU 12 is allocated to the real
address space as shown in FIG. 4, the local storage of

a partner VPU 12 can directly be accessed. This case however raises two large issues.

The first issue relates to a change in the VPU to which a VPU thread is dispatched. Assume that there are VPU threads A and B and they are executed by their respective VPUs 0 and 1 as shown in FIG. 32. Assume that the VPU threads A and B map the LSes (local storages) of their partner threads in their own EA spaces in order to cooperate with each other. Assume that LS0, LS1 and LS2 of VPU0, VPU1 and VPU2 are present in the RA space. In this case, it is the LS of a VPU executing the VPU thread B or the LS1 of the VPU1 that is mapped in the EA space of the VPU thread A. Conversely, it is the LS of a VPU executing the VPU thread A or the LS0 of the VPU0 that is mapped in the EA space of the VPU thread B. Assume that the scheduler of the VPU runtime environment changes a VPU to which the VPU thread A is dispatched and the VPU thread A is executed by the VPU 2. Since the VPU thread A is no longer executed by the VPU0, the LS of the VPU0, which is mapped in the EA space of the VPU thread B, becomes meaningless. In order to prevent the thread B from being aware of the change in the VPU to which the thread A is dispatched, the system needs to use some method for mapping the LS2 in the address of the EA space in which the LS0 is mapped and seeing the LS2 of the VPU2 through the thread B as the local

storage of the thread A.

The second issue relates to a correspondence between physical VPUs and logical VPUs. Actually, there are two levels to allocate VPUs to VPU threads.

5 The first level is to allocate logical VPUs to VPU threads and the second level is to allocate physical VPUs to the logical VPUs. The physical VPUs are real VPUs 12 managed by the virtual machine OS 301. The logical VPUs are virtual VPUs allocated to the guest
10 OSes by the virtual machine OS 301. This correspondence is also shown in FIG. 14. If the VPU runtime environment 401 manages the logical VPUs, the VPUs that are allocated to the VPU threads by the VPU runtime environment 401 are logical VPUs in FIG. 32.

15 FIG. 33 illustrates the concept of the above two level. The first issue corresponds to an issue of the assignment of VPU threads to logical VPUs in the upper stage in FIG. 33. The second issue corresponds to an issue of the allocation of physical VPUs to logical
20 VPUs in the lower stage in FIG. 33. In FIG. 33, three are selected from four physical VPUs and allocated to three logical VPUs, respectively. When a correspondence between the physical and logical VPUs changes, the setting needs to be changed appropriately
25 even though the allocation of logical VPUs to VPU threads does not change. For example, the entries of the page table corresponding to the local storages (LS)

have to be replaced to gain correct access to the LS of the changed logical VPU.

Assume that the virtual machine OS 301 allocates physical VPUs 1, 2 and 3 to their respective logical
5 VPUs 0, 1, 2 at a certain time, as shown in FIG. 34. In FIG. 34, the logical VPU1 is allocated to VPU thread A and logical VPU2 is allocated to VPU thread B. The VPU threads A and B map the LSes of the physical VPUs, which execute their partner threads, in their own EA
10 spaces. Specifically, LS3 of the physical VPU3, which executes the VPU thread B, is mapped in the EA space of the VPU thread A, and LS2 of the physical VPU2, which executes the VPU thread A, is mapped in the EA space of the VPU thread B. Assume that the virtual machine OS
15 301 allocates the physical VPUs 0 and 1 to the logical VPUs 0 and 1 again at a certain time. The physical VPU2, which is allocated to the logical VPU1 that executes the VPU thread A, is changed to the physical VPU1. The allocation of the logical VPUs to the VPU
20 threads does not change, but the correspondence between physical VPUs and logical VPUs changes. It is therefore necessary to change the LS of the physical VPU executing the VPU thread A, which is mapped in the EA space of the VPU thread B, from the LS2 of the
25 physical VPU2 to the LS1 of the physical VPU1 and gain correct access to the LS1 of the physical VPU1.

In order to resolve the above two issues described

above, the real-time processing system of the present embodiment controls the virtual memory mechanism such that the local storage of a VPU, which executes its partner thread, is always mapped in the fixed address of the EA space viewed from a thread. In other words, when the scheduler dispatches a logical VPU, or when the virtual machine OS changes a correspondence between physical and logical VPUs, the page table and segment table are rewritten appropriately to allow a thread executed by a VPU to see the local storage of a VPU that executes the partner thread at all times in the same address.

There now follows an explanation as to the relationship in EA space between two threads. The EA spaces of two threads are shared or unshared in the following three patterns:

1. Shared EA pattern: Two threads 1 and 2 share both the segment table and page table (FIG. 35).
2. Shared VA pattern: Two threads 1 and 2 share the page table and not the segment table but have their respective segment tables (FIG. 36).
3. Unshared pattern: Two threads 1 and 2 share neither the page table nor the segment table but have their respective page tables and segment tables (FIG. 37).

There now follows an explanation as to how the mapping of local storages of VPUs to the EA space are

controlled, taking the shared EA type as an example.

First, as shown in FIG. 38, address regions corresponding to the respective logical VPUs are arranged on the VA space. The contents of the page table are set up such that the local storages of physical VPUs corresponding to the logical VPUs are mapped to the address regions corresponding to the local storages of the logical VPUs. In this case, the local storages of the physical VPUs 0, 1 and 2 correspond to the address regions of the local storages of the logical VPUs 0, 1 and 2, respectively. Then, the segment table is set in such a manner that the thread A can see the local storage of a logical VPU that executes the thread B through segment a of a fixed address on the EA space. The segment table is also set in such a manner that the thread B can see the local storage of a logical VPU that executes the thread A through segment b of a fixed address on the EA space. In this case, the thread A is executed by the logical VPU2, and the thread B is executed by the logical VPU1. Assume here that the scheduler in the VPU runtime environment 401 dispatches the thread B to the logical VPU0. Then, the VPU runtime environment 401 automatically rewrites the segment table such that the thread A can see the local storage of the logical VPU0 that executes the thread B through the segment a, as shown in FIG. 39.

Assume here that a correspondence between the physical and logical VPUs changes because the virtual machine OS 301 dispatches the guest OS. As shown in FIG. 40, the VPU runtime environment 401 rewrites the page table such that the address regions of local storages of logical VPUs fixed on the VA space exactly correspond to the local storages of physical VPUs. In FIG. 40, since the physical VPUs 1, 2 and 3 change to the logical VPUs 0, 1 and 2, respectively, the page table is rewritten such that the address regions of local storages of the logical VPUs 0, 1 and 2 correspond to the local storages of the physical VPUs 1, 2 and 3.

As described above, when the logical VPU that executes a thread changes due to the dispatch of the thread, the segment table of mapping from EA space to VA space is rewritten to resolve the first issue. When a correspondence between physical and logical VPUs is changed by the virtual machine OS 301 or the like, the page table of mapping from VA space to RA space is rewritten to resolve the second issue.

The local memory (local storage) of a processor corresponding to the partner thread, which is mapped in the effective address space, is automatically changed in accordance with a processor that executes the partner thread. Thus, each thread can efficiently interact with its partner thread without being aware of

a processor to which the partner thread is dispatched. Consequently, a plurality of threads can be executed with efficiency and in parallel to one another.

5 The shared EA type has been described so far. In the shared VA type and unshared type, too, the first and second issues can be resolved by rewriting the segment table or the page table as in the shared EA type.

10 Another method of resolving the above first and second issues will be described taking the shared EA type as an example. If there are a plurality of VPU threads that run in cooperation with each other, the page table and segment table are set such that the local storages of VPUs that execute the threads are
15 consecutively mapped on the segment in the EA space. In FIG. 41, the thread A is executed by the physical VPU2 and the thread B is executed by the physical VPU0. The page table and segment table are set such that the local storages of the VPUs can consecutively be
20 arranged on the same segment. When the logical VPUs that execute the threads are changed by the scheduler in the VPU runtime environment 401 or the correspondence between physical and logical VPUs is changed by the virtual machine OS or the like, the page table
25 is rewritten to hide these changes from the threads A and B, and the mapping of VA and RA spaces is changed. FIG. 42 shows mapping in the case where the VPU that

executes the thread A is changed to the physical VPU1 and the VPU that executes the thread B is changed to the physical VPU3. Even though the changes are made, each of the threads A and B can always access the local
5 storage of the VPU that executes its partner thread by accessing a given area in the segment having a fixed address.

A procedure for address management performed by the VPU runtime environment 401 will now be described
10 with reference to the flowchart shown in FIG. 43. The VPU runtime environment 401 maps in the fixed address on the EA space of each thread an RA space corresponding to the local storage of the VPU that executes its partner thread (step S201). After that, the VPU
15 runtime environment 401 determines whether the VPU that executes the partner thread is changed due to a change in the VPU to which the partner thread is dispatched or a change in the correspondence between the logical and physical VPUs (step S202). If the VPU that executes
20 the partner thread is changed, the VPU runtime environment 401 rewrites the contents of the segment table or page table and changes the local storage mapped in the fixed address on the EA space of each thread in accordance with the VPU that executes the
25 partner thread (step S203).

The example described up to now is directed to a system for accessing a local storage of the VPU that

executes the partner thread. The system is suitable
for the tightly coupled threads that are always
executed simultaneously. However, there is a case
where the threads that run in cooperation with each
5 other are not always assigned to the VPU at once, as
in the loosely coupled thread group. In this case,
too, the EA space has a segment for mapping the local
storage of VPU 12 that executes the partner thread and
thus the segment is used as follows to deal with the
10 local storage.

First method: If a segment for mapping the local
storage of a VPU corresponding to a partner thread is
accessed while the partner thread is not running, a
thread is caused to wait until the partner thread
15 starts to run.

Second method: If a segment for mapping the local
storage of a VPU corresponding to a partner thread is
accessed while the partner thread is not running, a
thread becomes aware of it by an exception or an error
20 code.

Third method: When a thread exits, the contents
of the local storage, which are provided when the
thread runs finally, are stored in the memory area.
The mapping is controlled such that the entries of the
25 page table or segment table, which indicate the local
storage corresponding to the thread, indicate the
memory area. According to this method, even though the

partner thread is not running, a thread can continues to run as if there were a local storage corresponding to the partner thread. A specific example thereof is shown in FIGS. 44 and 45.

5 (1) Assume that threads A and B are executed by VPU0 and 1, respectively and the local storage LS0 of VPU0 that executes the thread A is mapped in the EA space of the thread B.

10 (2) When the thread A exits, the thread A or VPU runtime environment 401 stores (saves) the contents of local storage LS0 of VPU0 that executes the thread A in a memory area on the memory 14 (step S211).

15 (3) The VPU runtime environment 401 changes the address space for the local storage of the thread A, which is mapped in the EA space of the thread B, from the LS0 of VPU0 to the memory area on the memory 14 that stores the contents of the LS0 (step S212). Thus, the thread B can continue to run even after the thread A stops running.

20 (4) When a VPU is allocated to the thread A again, the VPU runtime environment 401 restores the content of the memory area on the memory 14 to the local storage of the VPU that executes the thread A (step S213). If the VPU0 is allocated to the thread A again, the content of the memory area is restored to
25 the local storage LS0 of the VPU0.

 (5) The VPU runtime environment 401 changes the

address space of the local storage of the thread A,
which is mapped in the EA space of the thread B, to the
local storage of the VPU that executes the thread A
(step S214). If the VPU0 is allocated to the thread A
5 again, the address space of the local storage of the
thread A, which is mapped in the EA space of the thread
B, is changed to the local storage LS0 of the VPU0.

If the VPU2 is allocated to the thread A, the
content of the memory area on the memory 14 is restored
10 to the local storage LS2 of the VPU2. Then, the
address space of the local storage of the thread A,
which is mapped in the EA space of the thread B, is
changed to the local storage LS2 of the VPU2.

State Transition of Threads

15 A thread generally makes a state transition from
when it is created until it is deleted. As shown in
FIG. 46, a thread makes the following seven state
transitions.

20 1. Not-existent state: This state is logical and
does not exist in an effective thread.

2. DORMANT state: A thread is created and does not
start running yet.

3. READY state: The thread is ready to start
running.

25 4. WAITING state: The thread waits for conditions
to meet to start (resume) running.

5. RUNNING state: The thread is actually running

on the VPU or MPU.

6. SUSPENDED state: The thread is forcibly suspended by the VPU runtime environment and other threads.

5 7. WAITING-SUSPENDED state: The waiting and suspended states overlap each other.

The conditions of transition between the above seven states and the thread contexts involved in the transition are as follows.

10 [Transition from NOT EXISTENT state to DORMANT state]

This transition is made by creating a thread.

A thread context is created but its contents are in the initial state.

[Transition from DORMANT state to NOT EXISTENT state]

15 This transition is made by deleting a thread.

If the thread is set to store its thread context, the stored thread context is discarded by the transition.

[Transition from DORMANT state to WAITING state]

20 This transition is made when the thread requests the runtime environment to schedule the thread.

[Transition from WAITING state to READY state]

25 This transition is made when an event (e.g., synchronization, communication, timer interruption) for which the thread waits is generated.

[Transition from READY state to RUNNING state]

This transition is made when the thread is

dispatched to MPU or VPU by the runtime environment.

The thread context is loaded. When the thread context is saved, it is restored.

[Transition from RUNNING state to READY state]

5 This transition is made when the running of the thread is preempted.

[Transition from RUNNING state to WAITING state]

 This transition is made when the thread suspends its own running to wait for an event using a
10 synchronization mechanism, a communication mechanism and the like.

 The thread in every class can be set to store its thread context. When a thread is set to store its thread context, the thread context is saved by the
15 runtime environment when the thread transits from RUNNING state to WAITING state. The saved thread context is maintained unless the thread transits to DORMANT state and restored when the thread transits to the RUNNING state.

20 [Transition from RUNNING state to SUSPENDED state]

 This transition is made when the running of the thread is forcibly suspended in response to an instruction from the runtime environment or other threads.

25 The thread in every class can be set to store its thread context. When a thread is set to store its thread context, the thread context is saved by the

runtime environment when the thread transits from
RUNNING state to SUSPENDED state. The saved thread
context is maintained unless the thread transits to
DORMANT state and restored when the thread transits to
5 the RUNNING state.

[Transition from RUNNING state to DORMANT state]

This transition is made when the thread in itself
exits its own running.

When the thread is set to store its thread
10 context, the contents of the thread context are
discarded by the transition.

[Transition from WAITING state to WAITING-SUSPENDED
state]

This transition is made when the thread is forced
15 to stop by instruction from outside while it is waiting
for an event to generate in the WAITING state.

[Transition from WAITING-SUSPENDED state to WAITING
state]

This transition is made when the thread resumes
20 running by instruction from outside while it is in the
WAITING-SUSPENDED state.

[Transition from WAITING-SUSPENDED state to SUSPENDED
state]

This transition is made when the event for which
25 the thread waits in the WAITING state is generated.

[Transition from SUSPENDED state to READY state]

This transition is made when the thread resumes

running by instruction from outside.

[Transition from READY state SUSPENDED state]

This transition is made when the thread stops running by external environment.

5 Execution Term of Thread

 The term of the running state of a thread to which a VPU is allocated is called an execution term. In general, a term from creation to deletion of a thread includes a plurality of execution terms of the thread.

10 FIG. 47 shows an example of thread states varied from creation to deletion. This example includes two execution terms during the presence of the thread. The thread context can be saved and restored using various methods. Most normal threads run so as to save a

15 context at the end of an execution term and restore the context at the beginning of the next execution term. In a certain periodic operation, the thread run so as to create a new context at the beginning of an execution term, use the context during the execution

20 term, and discard the context at the end of the execution term in every period.

Execution Term of Threads belonging to Tightly Coupled Thread Group

 FIG. 48 shows execution terms of threads belonging

25 to the same tightly coupled thread group. All the threads belonging to a certain tightly coupled thread group are scheduled by the VPU runtime environment 401

such that they can run at once in one execution term. This tightly coupled thread group is used chiefly for hard real-time threads. In order to achieve the operation, therefore, the VPU runtime environment 401
5 designates processors used at once and their number when an execution term is reserved for the hard real-time class. Moreover, the VPU runtime environment 401 makes contexts of threads running at once correspondent to the processors, respectively.

10 The threads, which belonged to the tightly coupled thread group in a certain execution term, can run separately from each other in other execution term by canceling their tightly coupled relationship. Each of the threads has to sense whether it runs as a tightly
15 coupled thread or separately from another thread and perform an operation of communication and synchronization with its partner thread. Each of the threads is provided with an attribute that indicates preemptive or non-preemptive. The preemptive attribute permits a
20 thread to be preempted during its execution term and, in other words, permits the thread to stop running. The non-preemptive attribute ensures that a thread cannot be preempted during its execution term. The non-preemptive attribute varies in meaning from thread
25 class to thread class. In the hard real-time class, when a thread starts to run, nothing but the thread in itself can stop the running until its execution term

ends. In the soft real-time class, preemptiveness is essential and thus the non-preemptive attribute is not supported. In the best effort class, a thread can be protected against being preempted from another best effort class, but it can be preempted from a higher-level class such as the hard real-time class and soft real-time class.

Execution Models of Threads

The execution models of threads can roughly be classified into two models: a periodic execution model as shown in FIG. 49 and an aperiodic execution model as shown in FIG. 50. In the periodic execution model, a thread is executed periodically. In the aperiodic running model, a thread is executed based on an event. The periodic execution model can be implemented using a software interrupt or an event object such as synchronization primitives. In the hard real-time class, the periodic execution model is implemented using a software interrupt. In other words, the VPU runtime environment 401 jumps to an entry point of a thread determined by a given method with timing to start a periodic operation or calls a callback function registered in advance by a given procedure. In the soft real-time class, the periodic execution model is implemented using an event object. In other words, since the VPU runtime environment 401 notifies a generation of a previously-registered event object in

each period, a soft real-time thread waits an event object in each period, and perform a given operation upon generation of the event, thereby realizing a periodic execution model. In the best effort class, the periodic execution model can be implemented using either one of a software interrupt or an event object. The actual execution does not always start at the beginning of each period, but may be delayed within constraints.

Using an event model, the aperiodic execution model can be realized as the periodic execution model. In the soft real-time class and best effort class, the aperiodic execution model differs from the periodic execution model only in the timing with which an event is notified and these models are the same in the implementing method. In the hard real-time class, the minimum inter-arrival time and the dead line, which are necessary for securing time requirements, strongly constrain the operation of the system; accordingly, the aperiodic execution is restricted.

Context Switching

In the real-time processing system according to the present embodiment, one of methods for switching a context at the end of the execution term of a VPU thread can be selected. Since the costs for switching the context are very high, the selection of one method improves the efficiency of switching. The selected

method is used at the end of the reserved execution term of a thread. When a context is switched during the execution term or at the time of preemption, all contexts of the current thread need to be saved in whatever case and restored when the thread resumes running next. For example, there are following methods of switching a VPU context.

1. Discard of Contexts

No contexts are saved.

2. Complete Saving of Contexts

All contexts of a VPU, including the states of the register and local storage of the VPU and those of the DMA controller in the memory controller, are saved.

3. Graceful Saving of Contexts

The context switching is delayed until all operations of the DMA controller in the memory controller in a VPU are completed. After that, the contents of the register and local storage in the VPU are saved. In this method, all the contexts of the VPU as well as the complete saving are saved.

One scheduler can be implemented to schedule both MPU and VPU threads and different schedulers can be done to schedule their respective MPU and VPU threads. Since the MPU and VPU threads differ in costs for switching a context, the implementation of different schedulers becomes more efficient.

Scheduling in Hard Real-Time Class

The scheduling of threads in the hard real-time class is performed using a reservation graph of an extended task graph. FIG. 51 shows an example of the task graph. The task graph represents a relationship between tasks. In FIG. 51, the arrows between tasks indicate the dependence of the tasks (relationship in input/output between the tasks). According to the example of FIG. 51, tasks 1 and 2 can freely start to run, a task 3 can start to run after both the tasks 1 and 2 stop running, and tasks 4 and 5 can start to run after the task 3 stops running. The task graph has no concepts of contexts. For example, when the tasks 1 and 4 should be processed using the same context, it cannot be described in the task graph. The following reservation graph of the extended task graph is therefore used in the real-time processing system of the present embodiment.

First, consider the task graph to be a relationship between not tasks but execution terms. By relating a context to each of the execution terms, a thread corresponding to the context runs in the execution term. If the same context is related to a plurality of execution terms, its corresponding thread runs in each of the execution terms. In the example shown in FIG. 52, the context of thread 1 is related to execution terms 1 and 2, and the thread 1 runs in each

of the execution terms 1 and 2. An attribute indicative of constraints of hard real-time ensured by the runtime environment is added to each of arrows between the execution terms in FIG. 52. Using a reservation graph so created, operation models and constraints such as time requirements of a real-time application can be described without making any modifications to the model of the real-time application. FIG. 53 shows an example of the reservation graph created based on the graph shown in FIG. 52. Contexts 1, 2 and 3 in FIG. 53 correspond to those of threads 1, 2 and 3 in FIG. 52, respectively.

Scheduling in Soft Real-Time Class

The scheduling of threads in the soft real-time class is performed using a fixed priority scheduling method in order to allow the running patterns of threads to be predicted. Two different scheduling algorithms are prepared for the scheduling method: one is fixed priority FIFO scheduling and the other is fixed priority round robin scheduling. In order to execute a higher-priority thread by priority, even while a lower-priority thread is running, the lower-priority thread is preempted and immediately the higher-priority thread starts to run. In order to avoid a priority inversion problem that occurs in a critical section, it is desirable to perform a synchronization mechanism such as a priority

inheritance protocol and a priority ceiling protocol.

Scheduling in Best Effort Class

The scheduling of threads in the best effort class is performed using dynamic priority scheduling and the like.

Hierarchical Scheduler

The scheduling function in the VPU runtime environment 401 can be fulfilled as a hierarchical scheduler as shown in FIG. 54. In other words, thread-level scheduling has two hierarchies of thread inter-class scheduling and thread intra-class scheduling. Thus, the scheduler in the VPU runtime environment 401 has a thread intra-class scheduling section 601 and a thread inter-class scheduling section 602. The thread inter-class scheduling section 602 schedules threads spreading over thread classes. The thread intra-class scheduling section 601 schedules threads belonging to each of thread classes. The section 601 includes a hard real-time (hard RT) class scheduling section 611, a soft real-time (soft RT) class scheduling section 612 and a best effort class scheduling section 613.

The thread inter-class scheduling and thread intra-class scheduling have a hierarchical structure. First, the thread inter-class scheduling operates to determine which thread class is executed and then which thread in the thread class is executed. The thread inter-class scheduling employs preemptive fixed

priority scheduling. The hard real-time class has the highest priority, with the soft real-time class and the best effort class following in that order. When a thread in a higher-priority class is ready to run, a lowest-priority thread is preempted. Synchronization between thread classes is achieved by a synchronous primitive provided by the VPU runtime environment 401. In particular, only the primitive can be used in a hard real-time thread to prevent a block from occurring in the hard real-time thread. When a best effort thread blocks a soft real-time thread, it is processed as a soft real-time thread to prevent priority from being inverted between thread classes. Furthermore, the use of, e.g., the priority inheritance protocol prevents another soft real-time thread from blocking the best effort thread.

Thread Parameters

In the real-time processing system according to the present embodiment, threads are scheduled using various parameters. The parameters common to the threads in each class are as follows:

- Class of threads (hard real-time, soft real-time, best effort);
- Resources for use (number of MPUs or VPUs, bandwidth, physical memory size, I/O device);
- Priority; and
- Preemptive or non-preemptive.

The following are parameters for the threads in the hard real-time class:

Execution term;

Dead line;

5 Period or minimum inter-arrival time; and

VPU context switching method.

FIG. 55 shows examples of fundamental parameters for the hard real-time class. In example 1 to designate an execution term shown in the uppermost part
10 of FIG. 55, one MPU and two VPUs are reserved at once in the designated execution term, and the context of each of the VPUs is completely saved. In this case, the threads run at the same time on the three processors and, after the execution term, the contexts
15 of VPU threads as well as that of an MPU thread are completely saved. In the upper right of FIG. 55, example 2 shows a method of designating a deadline to ensure that an operation represented by the number of VPUs and their execution term is performed before the
20 deadline. The deadline is designated by relative time starting at the request time when a reservation request is made. In the lowermost part of FIG. 55, example 3 shows a method of designating a periodic execution. In this example, an execution term that designates two
25 VPUs 12 is periodically repeated, and the contexts of VPU threads are discarded after the execution term for each period, with the result that all operations are

performed by new contexts. Moreover, the deadline is designated by relative time starting at the beginning of the period.

For example, there are following constraints as
5 other parameters used in the hard real-time class:

Timing constraints (absolute timing constraint and relative timing constraint);

Precedence constraint; and

Mutual exclusive constraint.

10 The timing constraints provide means for delaying execution timing. The absolute timing constraint is a condition for designating delay time with reference to static timing, such as the start time of a period, as shown in FIG. 56. The relative timing constraint is a
15 condition for designating permissible delay time with reference to dynamic timing and an event, such as the start time and end time of a certain, as shown in FIG. 57. Since the precedence constraint can be achieved by designating delay time as 0 or longer with
20 reference to the end time of a certain execution term using the relative timing constraint, it can be considered to be a special one for the relative timing constraint.

The mutual exclusive constraint is a condition for
25 ensuring that execution terms do not overlap each other, as shown in FIG. 58. The mutual exclusive constraint makes it possible to lessen the prediction

impossibility of the execution term, which is caused by a lock. In other words, all threads common to some resources are prevented from running at once to obviate a lock regarding the resources.

5 Synchronization mechanisms for Threads

 In the real-time processing system according to the present embodiment, the following synchronous primitives are used as synchronization mechanisms for threads:

10 Semaphore;
 Message queue;
 Message buffer;
 Event flag;
 Barrier; and
15 Mutex.

 The other synchronous primitives can be used. The real-time processing system of the present embodiment provides the following three methods to achieve the above synchronization mechanisms:

20 The synchronization mechanisms are implemented on the memory (main storage) 14 or the local storage 32 of a VPU using an instruction such as a TEST & SET;

 The synchronization mechanisms are implemented by hardware mechanisms such as a mail box and a signal
25 register; and

 The synchronization mechanisms are implemented using a mechanism provided as a service by the VPU

runtime environment.

Since the synchronization mechanisms have advantages and disadvantages, it is desirable to selectively use them according to the attributes of threads as shown in FIG. 59. In other words, a synchronization mechanism implemented using the memory (main storage MS) 14 that is shared and accessed by the MPU and VPUs can be used for threads in all classes. In contrast, a synchronization mechanism implemented on the local storage LS of a VPU 12 can be used only for threads belonging to the tightly coupled thread group. This is because only the threads belonging to the tightly coupled thread group ensure that their partner threads for synchronization run at the same. For example, if a thread belonging to the tightly coupled thread group is used for a synchronization mechanism implemented on the local storage of a VPU that executes the partner thread, the execution of the partner thread is ensured when the synchronization mechanism is used. Thus, the local storage of the VPU that executes the partner thread always stores information for the synchronization mechanism.

A synchronization mechanism using a means other than the memory (main storage MS) and local storage LS can be implemented by a hardware mechanism or a service of the VPU runtime environment 401. Since the threads belonging to the tightly coupled thread or those in the

hard real-time class require a high-speed synchronization mechanism, the synchronization mechanism implemented by the hardware mechanism is desirable to use in the threads. In contrast, the synchronization mechanism provided by the runtime environment is desirable to use in the threads belonging to the loosely coupled thread group or those belonging to the soft real-time class and best effort class.

Automatic Selection of Synchronization mechanism

In the real-time processing system according to the present embodiment, the above synchronization mechanisms can automatically be selected or switched in accordance with the attribute and status of threads. This operation is performed by a procedure as shown in FIG. 60. While threads for synchronization belong to the tightly coupled thread group (YES in step S201), a high-speed synchronization mechanism that is implemented by the memory 14, the local storage 32 of each VPU 12 or the hardware mechanism is used (steps S202, S203, S204, S205). When the threads change in status to cancel their tightly coupled relationship (NO in step S201), the high-speed synchronization mechanism is switched to a synchronization mechanism that is implemented as a synchronization mechanism on the memory 14 or a service of the VPU runtime environment 401 (steps S206, S207, S208).

The above switching can be provided for programs

running on the VPUs 12 in the form of a library or as a service of the VPU runtime environment 502 in each of the VPUs 12. A plurality of synchronization mechanisms can be switched as follows. The synchronization mechanisms can be secured in advance and used selectively or new synchronization mechanisms can be secured when the switching is performed.

For a synchronization mechanism using local storages of VPUs 12, threads needs to be executed at once by the VPUs like threads belonging to the tightly coupled thread group. This constraint is eased as follows. While a thread is not running, the contents of the local storage are stored in the memory 14 when the thread runs last, and mapping is so controlled that the stored contents are indicated by the entries of the page table or segment table indicating the local storage. According to this method, while the partner thread is not running, the thread can continue running as if there is a local storage related to the partner thread. When the thread starts to run by allocating a VPU 12 thereto, the contents stored in the memory 14 are restored to the local storage of the VPU 12 to change the mapping of a corresponding page table or segment table. Using a backup copy of the local storages of the VPUs 12, the synchronization mechanism using the local storages of VPUs 12 can be used even for threads that do not belong to the tightly coupled

thread group.

Reservation Graph

FIG. 61 shows a reservation graph corresponding to the data flow shown in FIG. 9. In FIG. 61, six boxes represent execution terms. The upper left number on each of the boxes indicates the ID of an execution term to be reserved. The symbol in each box indicates the identifier of a thread context related to the execution term. The lower right number on each box indicates the length (cost) of the execution term. The arrows connecting the boxes all denote precedence constraints. In other words, an arrow extending from one box to another box indicates that an operation in the execution term of the latter box starts after an operation in that of the former box is completed. The number with each arrow denotes an ID of a buffer used for data transfer between execution terms connected by the arrow, and the value with each number denotes the size of a buffer. The following are procedures 1 to 7 for performing operations in accordance with the reservation graph shown in FIG. 61.

1. Create a thread context that executes the DEMUX program 111 and call its identifier DEMUX.
2. Create a thread context that executes the A-DEC program 112 and call its identifier A-DEC.
3. Create a thread context that executes the V-DEC program 113 and call its identifier V-DEC.

4. Create a thread context that executes the TEXT program 114 and call its identifier TEXT.

5. Create a thread context that executes the PROG program 115 and call its identifier PROG.

5 6. Create a thread context that executes the BLEND program 116 and call its identifier BLEND.

7. Create a reservation request having a data structure as shown in FIG. 62 and sends it to the VPU runtime environment 401 to make a reservation.

10 According to each of the above procedures 1 to 6, if a program is designated to run as a thread, the VPU runtime environment 401 assigns necessary resources to the program to create a thread context. The handle of the thread context is returned and thus referred to as
15 an identifier.

FIG. 62 shows a reservation request containing buffer data written as BUFFER and execution term data written as TASK. The buffer data is used to declare a buffer on the memory 14 for data transfer between
20 execution terms. In the buffer data, "Id" indicates buffer number, "Size" indicates buffer size, "SrcTask" shows execution term number that writes data and "DstTask" shows execution term number that reads data. In the execution term data, "Id" represents execution
25 term number, "Class" indicates thread class (VPU shows VPU thread and HRT shows hard real-time class. In addition to these, there are MPU showing MPU thread,

SRT showing soft real-time class, BST showing best effort class and so on), "ThreadContext" denotes thread context corresponding to the execution term, "Cost" indicates length or cost of the execution term,

5 "Constraint" represents various constraints based on the execution term, "InputBuffer" shows a list of identifiers of buffers read in the execution term and "OutputBuffer" indicates a list of identifiers of buffers written in the execution term. The

10 "Constraint" also can include "Precedence" showing precedence constraint, "Absolute Timing" showing absolute timing constraint, "Relative Timing" showing relative timing constraint and "Exclusive" showing mutual exclusive constraint. The "Constraint" has a

15 list of numbers of execution terms of partner threads for constraints.

The buffer area reserved by the reservation request shown in FIG. 62 is allocated to the main memory 14 and released therefrom by the VPU runtime

20 environment 401. The allocation of the buffer area is performed when a thread that writes data to the buffer area starts to run. The release of the buffer area is performed when a thread that reads data from the buffer area exits. The thread can be notified of the address

25 of the allocated buffer using an address, a variable or a register that is predetermined when the thread starts to run. In the real-time processing system of the

present embodiment, when the program module 100 shown in FIG. 7 is provided, the structural description 117 shown in FIG. 8 is read out of the program module 100 and, based on the structural description 117, a thread
5 context is created by the above procedures and a reservation request as shown in FIG. 62 is created and issued, thereby providing a function of executing the program module 100. This function allows the operation of dedicated hardware described by the program module
10 100 as shown in FIG. 7 to be performed by processing software by a plurality of processors. A program module having a structure as shown in FIG. 7 is created for each hardware to be implemented and then executed by an apparatus having a function conforming to the
15 real-time processing system of the present embodiment, with the result that the apparatus can be operated as desired hardware.

Providing the reservation request shown in FIG. 62, the VPU runtime environment 401 determines
20 which VPU 12 executes each task with which timing in a period. This is scheduling. Actually, a plurality of reservation requests can be provided at once; therefore, operation timing is determined to prevent them from contradicting each other (prevent given
25 constraints from not being satisfied). Assuming that only the reservation request shown in FIG. 62 is made when there are two VPUs 12 as shown in FIG. 63,

the scheduling is performed such that the VPU 0 sequentially performs DEMUX, V-DEC, PROG and BLEND operations which cannot be done in parallel and after the DEMUX operation, the VPU1 performs the A-DEC and
5 TEXT operations that can be done in parallel.

Software Pipeline

If there is no time enough to perform the DEMUX, V-DEC, PROG and BLEND operations in sequence within one period, software pipeline processing is carried out
10 over a plurality of periods. For example, as shown in FIG. 64, the VPU 0 performs the DEMUX and V-DEC operations in the first period and the VPU 1 performs the A-DEC, TEXT, PROG and BLEND operations in the second period. In the second period, the VPU 0
15 performs DEMUX and V-DEC operations in the next frame in parallel with the A-DEC, TEXT, PROG and BLEND operations. In other words, as shown in FIG. 65, the pipeline processing is performed in which the VPU 1 performs the A-DEC, TEXT, PROG and BLEND operations
20 upon receipt of outputs from the DEMUX and V-DEC operations in the preceding period while the VPU 0 is performing the DEMUX and V-DEC operations.

Scheduling in Consideration of the Number of Buffers

When a buffer is used to transfer data between a
25 thread running in an execution term and a thread running in another execution term, the buffer is occupied from the beginning of the execution term on

the data write side to the end of the execution term on
the data read side. For example, as shown in FIG. 66,
when a buffer on the memory 14 (main storage) is used
to transfer data between execution terms A and B, it is
5 occupied from the beginning of execution term A to the
end of execution term B. Therefore, when a buffer is
used to transfer data from execution term A to
execution term B and the execution terms A and B belong
to their respective periods adjacent to each other in
10 software pipeline processing, the number of buffers
required varies according to the execution timing in
the execution terms A and B. For example, as shown in
FIG. 67, when threads are scheduled such that they run
in the execution term A earlier than in the execution
15 term B in each period, data is transferred from
execution term A_n (A_n means execution term A in period
 n) to execution term B_n in the next period, and data is
transferred from execution term A_{n+1} to execution term
 B_{n+1} in the next period. Since the execution term A_{n+1}
20 is interposed between A_n and B_n , the buffer for
transferring data from A_n to B_n cannot be used for
transferring data from A_{n+1} to B_{n+1} but a new buffer
has to be used. In other words, two buffers are
required. On the other hand, as shown in FIG. 68, when
25 threads are scheduled such that they start to run in
execution term A after the end of execution term B in
one period, data that is written to a buffer in

execution term A_n is read out of the buffer in
execution term B_n . Then, data is written to the same
buffer in execution term A_{n+1} and read therefrom in
execution term B_{n+1} . That is, a single buffer has only
5 to be used.

In the real-time processing system according to
the present embodiment, the scheduler in the VPU
runtime environment 401 schedules execution terms to be
reserved such that the amount of use of buffer memory
10 areas becomes as small as possible. More specifically,
in order to execute software pipeline processing of
two VPUs 0 and 1, the scheduler in the VPU runtime
environment 401 divides an operation into two partial
operations (one to be performed first by the VPU 0 and
15 the other to be performed next by the VPU 1) as shown
in the flowchart in FIG. 69 (step S211). Then, the
scheduler extracts threads (thread A in the partial
operation to be performed first and thread B in the
partial operation to be performed next) which
20 inputs/outputs data through a buffer between the two
VPUs (step S212). The threads A and B are scheduled
such that the thread A starts to run after the end of
the execution term for the thread B in each period
(step S213).

25 Reservation Graph having a Hierarchical Structure

Though the reservation graph shown in FIG. 61 has
no hierarchical structure, a reservation graph having a

hierarchical structure can be used as shown in FIG. 70. In FIG. 70, the execution term A precedes the execution term B and the execution term B precedes the execution term C. In the execution term B, the execution term D precedes execution terms E and F. Resolving the hierarchy, the execution term A precedes the execution term D and the execution terms E and F precede the execution term C.

Reservation Request made in Consideration of Tightly Coupled Thread Group

In the reservation graph shown in FIG. 61, when a thread executing the V-DEC and a thread executing the PROG belong to the tightly coupled thread group, a reservation request indicative of the coupled attribute is created as shown in FIG. 71. In this reservation request, "TightlyCoupled" indicates an ID of the execution term corresponding to the partner thread.

The above threads are therefore scheduled as shown in FIG. 72 such that they can be executed at once by different VPUs. In this case, the threads can communicate with each other via a local storage and thus no buffers need to be provided on the memory 14.

Scheduling Algorithm based on Structural Description

There now follows descriptions as to a procedure for reserving an execution term of each thread based on the structural description incorporated into the program module.

FIG. 8 shows an example of the structural description 117 incorporated in the program module 100 shown in FIG. 7. With the structural description 117, the VPU runtime environment 401 performs the following steps.

1. The programs that are written in the module field of the structural description 117 are loaded to generate threads that execute the programs.

In the present embodiment, one thread is generated for each of entries of the structural description 117. If the structural description 117 includes entries having the same module name, a plurality of threads that execute the same module are generated so as to correspond to their respective entries. In the example of FIG. 8, all threads are generated to belong to one process; however, the threads can belong to different processes or thread groups can belong to different processes.

2. A reservation request having a data structure as shown in FIG. 62 is created based on the information of the structural description 117.

3. The reservation request is sent to the VPU runtime environment to schedule the threads and start to run the threads.

The above step 2 of creating the reservation request is performed as follows.

First, BUFFER records are created to correspond to

the output fields of the structural description 117 in a one-to-one basis and added to the reservation request. For instance, in the example of FIG. 8, the second output data of the DEMUX module is supplied to the V-DEC through the 1-MB buffer, so that a BUFFER record whose Id is 2 as shown in FIG. 62 is created. In this BUFFER record, the buffer size is described as 1MB in Size field, reference to TASK record whose Id is 1 and which corresponds to a DEMUX module that writes data to the buffer is described in SrcTask field, and reference to TASK record whose Id is 3 and which corresponds to a V-DEC module that reads data from the buffer is described in DstTask field.

Then, TASK records are created to correspond to the module fields of the structural description 117 on a one-to-one basis and added to the reservation request. For instance, in the example of FIG. 8, a TASK record whose Id is 3 as shown in FIG. 62 is created as one corresponding to the V-DEC module. This TASK record has the following information.

Class field: Flag to indicate what attribute is used to execute a thread designated in the TASK record.

In this field, "VPU" represents a thread that runs on the VPU and "HRT" shows a thread in the hard-real time class. These information items are set based on the information described in the thread parameters of the structural description 117 shown in FIG. 8.

ThreadContext field: Flag to designate a thread context of a thread whose running is to be reserved in the TASK record. More specifically, a program module designated in the module field of the structural description 117 is loaded, a thread that executes the program module is generated by the VPU runtime environment 401, and an identifier (a pointer or the like) of the thread context of the thread is recorded in the "ThreadContext" field.

10 Constraint field: Flag to record constraints of the TASK record. When the constraint is precedence constraint, a required number of Ids of another TASK record preceded by the TASK record is designated after the "Precede" field. For example, a TASK record whose
15 Id is 3 precedes a TASK record corresponding to the PROG module whose Id is 5.

 InputBuffer field: Flag to designate a required number of Ids of the Buffer record of a buffer from which data is read by the thread designated by the TASK
20 record.

 OutputBuffer field: Flag to designate a required number of Ids of the Buffer record of a buffer to which data is written by the thread designated by the TASK record.

25 If the structural description is provided as discussed above, its corresponding reservation request is created.

When the reservation request is sent to the scheduler in the VPU runtime environment 401, the scheduler creates a schedule necessary for performing the reservation request. This schedule represents
5 which VPU is allocated to which thread with which timing and how long the VPU is allocated in a period as shown in FIG. 63. Actually, the schedule can be represented by a reservation list as shown in FIG. 73.

The reservation list shown in FIG. 73 includes
10 reservation entries related to the respective VPUs. Each of the reservation entries includes a start time field indicating when a thread is executed by VPU in each period (execution start timing of the thread), an execution term field indicating how long the VPU is
15 allocated to the thread (execution term of the thread), and a running thread field indicating an identifier of the thread. The reservation entries are sorted in order of start time according to the VPUs and linked to the reservation list.

20 The procedure for creating a reservation list as shown in FIG. 73 from the reservation request shown in FIG. 62 or FIG. 71 can be carried out by the flowchart shown in FIG. 74.

Basically, the TASK records in the reservation
25 request have only to be sequenced in consideration of the relationship in input/output using BUFFER and the running time of VPUs has only to be assigned to each of

the TASK records in the order of data flow. It is then necessary to simultaneously allocate the VPUs to the TASKs belonging to the tightly coupled thread group.

5 The procedure is shown in FIG. 74. Upon receiving a reservation request, the VPU runtime environment 401 schedules all the tasks designated by TASK records in the reservation request by the following steps (in other words, the VPU runtime environment 401 creates a reservation list for reserving a VPU to which each task
10 is assigned and the execution start timing and execution term of the task).

 Step S301: The VPU runtime environment 401 selects a task whose all of preceding tasks (input tasks) have been already scheduled, and which have no tightly
15 coupled attributes, from among tasks that are not scheduled. If a task is preceded by no input tasks, it is determined as one whose input tasks have been already scheduled.

 If there is a task whose input tasks have been
20 already scheduled, and which have no tightly coupled attributes, the VPU runtime environment 401 selects it and moves to step S302. If not, it moves to step S304.

 Step S302: If there is a VPU that can assign the execution start timing and execution term of the
25 selected task under satisfactory constraints, the VPU runtime environment 401 moves to step S303. If not, the VPU runtime environment 401 fails in the scheduling

and makes a notification of the fail.

Step S303: The VPU runtime environment 401 creates reservation entries of the selected task and links them to the reservation list.

5 Step S304: The VPU runtime environment 401 selects tasks whose all input tasks have been already scheduled, and that belong to a tightly coupled group, from among tasks that are not scheduled. If tasks are preceded by no input tasks, they are determined as ones
10 whose input tasks have been already scheduled.

 If there are tasks whose input tasks have been already scheduled, and which belong to the tightly coupled group, the VPU runtime environment 401 selects them and moves to step S305. If not, it ends
15 scheduling.

 Step S305: If there are VPUs that can reserve all tasks included in the selected tasks at once (to have the same execution start timing and the same execution term), the VPU runtime environment 401 moves to step
20 S306. If not, the VPU runtime environment 401 fails in the scheduling and makes a notification of the fail.

 Step S306: Reservation entries of all tasks of the selected set of tasks are created and linked to the reservation list.

25 The steps of scheduling for one reservation request has been described. Actually, a plurality of reservation requests are usually present at once in one

system. In this case, the reservation requests can be scheduled through the above steps and, more favorably, they can be done simultaneously through the above steps.

5 The present embodiment has been described taking the program module describing the operations of a digital TV broadcast receiver as an example. If, however, a program module describing the operations of various types of hardware is prepared, the operations
10 of hardware can be performed by software.

 The MPU 11 and VPUs 12 provided in the computer system shown in FIG. 1 can be implemented as parallel processors mixed on one chip. In this case, too, the VPU runtime environment executed by the MPU 11 or the
15 VPU runtime environment executed by a specific VPU or the like can control scheduling for the VPUs 12.

 If the programs running as the VPU runtime environment or the programs of the operating system including the VPU runtime environment are stored in a
20 computer readable storage medium and then introduced and executed in a computer including a plurality of processors each having a local memory, the same advantages as those of the foregoing embodiment of the present invention can be obtained.

25 Additional advantages and modifications will readily occur to those skilled in the art. Therefore, the invention in its broader aspects is not limited to

the specific details and representative embodiments
shown and described herein. Accordingly, various
modifications may be made without departing from the
spirit or scope of the general inventive concept as
5 defined by the appended claims and their equivalents.